arXiv:1611.09072v2 [cs.DS] 30 Nov 2016

FASTER WORST-CASE UPDATE TIME FOR DYNAMIC SUBGRAPH CONNECTIVITY

RAN DUAN AND LE ZHANG

Abstract. The *dynamic subgraph connectivity* problem asks how to maintain an understanding of connectivity under vertex updates. An update may switch "on" or "off" a vertex, and might insert a new edge or delete an existing edge. Queries test the connectivity in the subgraph induced by active ones. Former results on this problem are data structures supporting vertex update in $\widetilde{O}(m^{2/3})$ amortized time, and for worst case the latest result are $\widetilde{O}(m^{4/5})$ and $\widetilde{O}(\sqrt{mn})$ [6, 10, 2]. In this paper, we give a structure of linear space, with worst-case update time $\widetilde{O}(m^{3/4})$. The data structure is randomized, as a Monte Carlo dynamic graph algorithm [24] is used as a subroutine.

1. Introduction

In this paper we consider the following problem: Given a graph $G = (V, E)$, edges in E may be deleted, and new edges might be inserted into E. Meanwhile, there is a subset S of V, which contains active vertices in V. S evolves over time, as a vertex status may switch between *on* and *off* by a vertex update. Given $u, v \in S$, a query tests whether they are connected in the subgraph of G induced by S.

This problem is called *dynamic subgraph connectivity*, which was first proposed by Frigioni and Italiano [16]. It captures the whole set of possible updates in a dynamic graph. In a real communication network, besides link updates (which are cables established or failed), vertex updates (which correspond to processors failure and reboot) are conceivable to occur more frequently. In [16], this model is referred to as *complete dynamic model*, and poly-logarithmic connectivity results were described for the special case of planar graphs. For general graphs, it was first considered by Chan [5]. For a graph with m edges, the method proposed supports $\widetilde{O}(m^{4\omega/(3\omega+3)})$[1] amortized update time, which is $O(m^{0.94})$ following from FMM (*Fast Matrix Multiplication*) algorithm of [7], where ω is the matrix multiplication exponent. Its space usage is linear, and has query time of $\widetilde{O}(m^{1/3})$. Later, Chan, Pătraşcu, and Roditty [6] showed a simpler algorithm, having an improved bound of $\widetilde{O}(m^{2/3})$. The new algorithm is of compact description, getting rid of the use of FMM. Space usage of it has increased to $\widetilde{O}(m^{4/3})$, however. With the same update time, Duan [10] gave another structure occupying linear space. Also in [10], algorithms with worst-case sublinear update time bounds were described. Its worst-case update time is $\widetilde{O}(m^{4/5})$, and uses as a subroutine $O(\sqrt{n})$ worst-case deterministic dynamic connectivity structure, which is Frederickson's $O(\sqrt{m})$ worst-case algorithm [15] sped up using *sparsification* technique proposed by Eppstein et al. [14].

[1] $\widetilde{O}(\cdot)$ hides poly-logarithmic factors.

Dynamic subgraph connectivity is a generalization of dynamic graph connectivity, a well-studied and fundamental problem in the area of dynamic graph. The latter confines to edge updates. For amortized circumstances, Henzinger and King [20] were the first to propose a randomized Las Vegas algorithm with poly-logarithmic update time. And it was then improved by Holm et al. [21] to a deterministic solution with $O(\lg^2 n)^2$ time per update. And again, Thorup [30] gave a randomized Las Vegas algorithm with $O(\lg n(\lg\lg n)^3)$ time bound, which is recently improved to $O(\lg n(\lg\lg n)^2)$ by Huang et al. [22]. The latest result for deterministic amortized time is by [32], with update time being $O(\lg^2 n/\lg\lg n)$. For worst-case scenario, it had been no improvement for a very long time. As a significant breakthrough, Kapron et al. [24] presented a randomized Monte Carlo algorithm with worst-case poly-logarithmic update time. Recently, it has several improvements [17, 31]. Also for dynamic graph connectivity problem, there is a cell probe lower bound of $\Omega(\lg n)$ per operation proved by Pătraşcu and Demaine [27].

Other related problems on subgraph connectivity, especially on vertex failures, are distance oracles avoiding a single node or link failure [3, 4, 9], and dual-failure distance oracles [11]. Also in [13, 12, 19], connectivity oracles for failure-prone graphs were proposed. Rather than a single vertex, this structure handles d failed vertices, and can process in $\widetilde{O}(d^3)$ time, then answers whether u and v are connected in $O(d)$ time. It is a counterpart of edge-failure oracle, which was considered in [28]. There an algorithm capable of recovering from a set of d edge updates in time $O(d \operatorname{poly}\lg n)$ was given. Besides, there are many applications of subgraph connectivity to geometric connectivity. See [5, 6] for details.

1.1. **Our Results.** With the new advances in worst-case edge-update model of dynamic graph connectivity [24, 17, 31], we consider the problem of whether poly-logarithmic worst-case edge-update time can bring progress in vertex-update model. The answer is affirmative. However, it does not come automatically by replacing $O(\sqrt{n})$ deterministic result [15, 14] with poly-logarithmic randomized one in [10]. It also cannot be achieved by tuning parameters of the former result. As an intuitive illustration, one cannot get a worst-case update time of $\widetilde{O}(m^{2/3})$ by simply replacing the *amortized* poly-logarithmic edge-update structure used in [6] with the *worst-case* poly-logarithmic one.

In this paper, we present a new algorithm with worst-case update time $\widetilde{O}(m^{3/4})$, query time $\widetilde{O}(m^{1/4})$. Answers to queries are correct w.h.p., with one-sided error. It occupies linear space, and is as simple as amortized $\widetilde{O}(m^{2/3})$ structures proposed in [6, 10]. The novelty is a new organization of former auxiliary data structures. Vertices are classified according to degrees. Components of the subgraph of low-degree vertices are further partitioned into two types. And connectivity of vertices of higher degree via these components is captured in a subtle way. Auxiliary data structures should be combined nicely so that properties of the whole graph can emerge. Note that it is far from being trivial to get the update time improved. This is at least partially demonstrated by the fact that we only get rid of an $m^{0.05}$ factor, whereas now edge-update time is *poly-logarithmic* rather than *polynomial*, not showing any sign of being "proportional."

Recent results on dynamic subgraph connectivity are summarized in Table 1. The second result comes from an application of dynamic DFS tree [2], which is

[2]We use $\lg x$ to denote $\log_2 x$.

TABLE 1. Results on Dynamic Subgraph Connectivity

Update time	Query time	Notes
$\widetilde{O}(m^{4\omega/(3\omega+3)})$	$\widetilde{O}(m^{1/3})$	Amortized, deterministic, linear space [5]
$\widetilde{O}(\sqrt{mn})$	$O(1)$	Worst case, deterministic, space $\widetilde{O}(m)$ [2]
$\widetilde{O}(m^{2/3})$	$\widetilde{O}(m^{1/3})$	Amortized, deterministic, space $\widetilde{O}(m^{4/3})$ [6]
$\widetilde{O}(m^{2/3})$	$\widetilde{O}(m^{1/3})$	Amortized, deterministic, linear space [10]
$\widetilde{O}(m^{4/5})$	$\widetilde{O}(m^{1/5})$	Worst case, deterministic, space $\widetilde{O}(m)$ [10]
$\widetilde{O}(m^{3/4})$	$\widetilde{O}(m^{1/4})$	Worst case, randomized, linear space, this paper

again a periodic rebuilding technique with fault tolerant DFS trees (see [2] for the definitions of related concepts). On one side, with Kapron et al.'s result [24] , we can just maintain a dynamic spanning forest to get an update time of $\widetilde{O}(n)$. With such a high update time, $O(1)$ query time follows by explicitly updating the information about which component a vertex belongs to after every update. Although the resulted solution is Monte Carlo, without loss of generality, assuming $m = \Omega(n)$, $\sqrt{mn}$ is always $\Omega(n)$. On the other side, since $n = \Omega(m^{1/2})$, $\sqrt{mn}$ is always $\Omega(m^{3/4})$, which indicates that our update time is always no worse than $\widetilde{O}(\sqrt{mn})$.

For all the bottom four results listed in Table 1, we can trade faster query time with slower update time.[3] For our result, you can have a version of $\widetilde{O}(m^{3/4+\epsilon})$ update and $\widetilde{O}(m^{1/4-\epsilon})$ query time. Details are shown in Appendix A. And in particular, a data structure of $\widetilde{O}(m^{4/5})$ update and $\widetilde{O}(m^{1/5})$ query time can be achieved as [10], by extending the result in this paper, although randomized. But from [10], one cannot get an algorithm with update time being $\widetilde{O}(m^{3/4})$. The trade-off phenomenon is definitely hard to break, as indicated in Section 5.

Our result is grouped as the following theorem.

Theorem 1.1. *Given a graph $G = (V, E)$, the data structure for dynamic subgraph connectivity has worst-case vertex (edge) update time $\widetilde{O}(m^{3/4})$, query time $\widetilde{O}(m^{1/4})$, where m is the value of $|E|$ at the time of the operation. The answer to each query is correct if the answer is "yes", and is correct w.h.p. if the answer is "no." Preprocessing time is $\widetilde{O}(m^{5/4})$, and space usage is linear.*

The paper is organized as follows. Section 2 specifies relevant notations. Our algorithm is given in Section 3. After this, its extension is discussed in Section 4. And final remarks are made in Section 5.

[3]The trade-off issues for amortized one had not been explicitly discussed in [10]. But we observed that it can be extended.

2. Notation and tree representation

A graph is represented by $G = (V, E)$, with $n = |V|$, $m = |E|$. Elements in V are referred to as *vertices* of G, whereas E is the set of *edges*. When G is a tree T, $v \in V$ is also named as a *node* of T.

2.1. **Trees as Euler tours.** An *Euler tour* (*cycle*) for a digraph is a cycle that traverses each edge exactly once, although it may visit a vertex more than once. For a strongly connected digraph, it has an Euler tour iff *in* degree and *out* degree are equal for each vertex. We will adopt Euler tours to represent dynamic trees. This representation was first proposed by Miltersen et al. [26] and independently by Henzinger and King [20].

Given a tree T, following the idea of Tarjan's [29], we replace each edge $\{v, w\}$ by two directed arcs (v, w) and (w, v), and add a self-loop (v, v) for each node v. The resulted digraph always has at least one Euler tour. Choose one such tour, and break it at an arbitrary place. And now we get a list of directed arcs. The list can be further stored as a red-black tree, which is called *Euler-tour* tree, a.k.a. *ET*-tree.

The list for T gives an Euler-tour order of nodes of T. Recall the self-loop (v, v) for node v in ET-tree. We can extract the order of these self-loops, i.e. nodes, and represent the order in another red-black tree. We call it the *order* tree, for nodes of T.

Observation 2.1. Given two nodes in T, the order of them w.r.t. Euler-tour order can be determined by order tree in $O(\lg n)$ worst-case time.

As shown in [29], ET-tree supports $O(\lg n)$ worst-case time for *linking* and *cutting*. We show briefly here how to get order trees for two resulted subtrees after cutting. Suppose the update is $cut(\{v, w\})$, and the list for T is like this, $\langle L_1, (v, w), L_2, (w, v), L_3\rangle$. Then lists for the two resulted subtrees should be $\langle L_1, L_3\rangle$ and $\langle L_2\rangle$. Augment ET-tree to enable us to find the first and last self-loop edges, i.e. nodes, in $O(\lg n)$ time. (This can be achieved by a boolean bit at each node of ET-tree indicating whether there exists a self-loop in the subtree rooted at the red-black tree node.) We query ET-tree($\langle L_2\rangle$) to get the first and last nodes, a and b. And order trees for two resulted subtrees can be obtained by splitting the initial order tree before a and again after b, and joining the first and last split parts.

Observation 2.2. ET-trees and order trees can be updated in $O(\lg n)$ worst-case time subjected to a link or cut operation on dynamic trees.

3. The data structure

Given a graph $G = (V, E)$, for every $v \in V$, a vertex v' and an edge (v, v') are added. A trivial observation follows that u, v are connected in G iff they are connected in the variant, which has $m = \Omega(n)$. It is required in the definition of dynamic subgraph connectivity problem [5]. It facilitates our presentation of time and space complexity as functions of m, rather than n in the case of degenerate graphs; hereafter $m = \Omega(n)$ is implicit for G. We assume G is static. The case of edge updates of E is dealt with in Section 4.

Main ideas are to remove vertices of low degrees. Artificial edges are added to capture the loss of connectivity due to the removal. As G evolves dynamically, artificial edges update accordingly. The subtle point is how to maintain these

artificial edges consistently and efficiently. We partition V based on their degrees w.r.t. G. Use $\deg_G(v)$ to denote the degree of v in G.

- C : Vertices with $\deg_G(v) > m^{1/2}$
- B: Vertices with $m^{1/4} < \deg_G(v) \leq m^{1/2}$
- A: Vertices with $\deg_G(v) \leq m^{1/4}$

Recall that S is the set of active vertices. Further, we select vertices $C \cap S$, $B \cap S$, and $A \cap S$ as V_C, V_B, and V_A respectively. Our goal is to remove vertices in V_A.

Consider the subgraph G_A of G induced by V_A. Define *total degree* of a component as the sum of $\deg_G(v)$'s for v's in it. We again partition components of G_A into two types: *high* component, with total degree $> m^{1/4}$; *low* component, with total degree $\leq m^{1/4}$. And for G_A, we also maintain a spanning forest F of it by a deterministic dynamic connectivity data structure proposed in [25]. Conceptually, it is simpler than Frederickson's approach. And update time is also faster.

Theorem 3.1 ([25])**.** *A spanning forest F of G_A can be maintained by a deterministic dynamic connectivity data structure demanding linear space, with $O(\sqrt{m(\lg\lg n)^2/\lg n})$ worst-case time for edge update in G_A, and constant query time to determine whether two vertices are connected in G_A.*

Note that a data structure with $\widetilde{O}(m)$ update time suffices. Since vertices in V_A are of degree $\leq m^{1/4}$, update time of a vertex switch for G_A is $\widetilde{O}(m^{3/4})$.

3.1. **Path graph.** Recall that the goal is to remove V_A. Path-graph structure adds artificial edges to restore "being-connected" relation of some vertices in V_B linking to the same component of G_A. It was first proposed in [10]. The crucial point is its efficiency. It is capable of being updated in $\widetilde{O}(1)$ time to reflect an *edge* update in G_A or involving a vertex in V_B.

W.l.o.g., assume that $V = \{0, 1, \ldots, n-1\}$. And now focus on V_B and a spanning tree T of G_A.

- *subpath* tree: For every node v in T, we can determine vertices in V_B adjacent to v. Store these vertices of V_B in a red-black tree, ordered by their id's. And call it subpath tree of v.

- *path* tree: Some nodes in T have adjacent vertices in V_B, whereas others do not. Anyway, we group former nodes, i.e. nodes having adjacent vertices in V_B, as a red-black tree according to Euler-tour order of T. Call it path tree of T.

For every node v in T, we build a *subpath* linking vertices in subpath tree of v in order. Now, for T, we concatenate subpaths associating with path-tree nodes of T, i.e. nodes of T which have adjacent vertices in V_B, in order as a *path*. The resulted path is called *path graph* for V_B w.r.t. T. An example is shown in Fig. 1.

Next we discuss some properties of path graphs. It should be updated subjected to vertex updates in V_B, and edge updates in F. (Recall that F is a spanning forest maintained for G_A.)

- Reflect vertex updates in V_B:
  For clarity, only "off" case is covered here. Suppose the vertex is $v \in V_B$. It has at most $m^{1/2}$ edges linked with V_A.

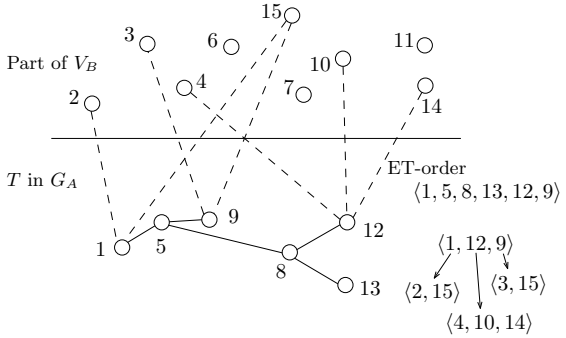

FIGURE 1. Path graph for V_B w.r.t. a spanning tree T in G_A. Edges within part of V_B shown here are omitted for clarity, and *dashed* edges emphasize the fact of these edges being between V_B and V_A. The path tree is on sequence $\langle 1, 12, 9\rangle$, and three subpath trees are on sequences $\langle 2, 15\rangle$, $\langle 4, 10, 14\rangle$, and $\langle 3, 15\rangle$ respectively. The resulted path graph is a path $\langle 2, 15, 4, 10, 14, 3, 15\rangle$.

Consider such an edge (v, w), where w is a node of T. We find w in path tree of T first. Then we know the subpath associated with w. Next just remove the two edges connecting v to its previous vertex, and to its next vertex in the subpath. And add an edge to relink previous and next vertices. If v is the 1st vertex of the subpath, we need to find the predecessor of w in path tree, and do some removal and relinking work. Similar operations work for v being the last.

Total cost for it is $\widetilde{O}(m^{1/2})$.

- Reflect edge updates in F:

  Take cutting as an example. As before, we assume the list for T is $\langle L_1, (v, w), L_2, (w, v), L_3\rangle$, and the resulted two subtrees are $\langle L_1, L_3\rangle$ and $\langle L_2\rangle$. And also suppose the first and last self-loop edges, i.e. nodes, of $\langle L_2\rangle$ are a and b. With the help of order tree for T, we can determine the predecessor of a in path tree in $O(\lg^2 n)$ time. Also the same cost suffices for finding the successor of b in path tree. After getting predecessor and successor, splitting path tree can be done. We update $O(1)$ edges in path graph to reflect splitting and joining of path trees.

  As a conclusion, path graph can be updated in $\widetilde{O}(1)$ time to reflect an edge update among F.

We group the results above into a lemma.

Lemma 3.2. *Path graphs can be maintained in $\widetilde{O}(m^{1/2})$ time for a vertex switch in V_B, and in $\widetilde{O}(1)$ time for a link or cut operation among F.*

3.2. **Adjacency structure.** This structure was first discussed in [10]. With it, given a component of G_A and a vertex in C, a query on whether these two are connected can be answered in $\widetilde{O}(1)$ time. It is not a static structure. It is updated synchronously while G evolves.

Focus on a vertex v in C and F.

- *sub-adjacency* tree: For a spanning tree T in F, several nodes of it might be adjacent to v. We group these nodes of T as a red-black tree, ordered by Euler-tour order of T. And the search tree is called sub-adjacent tree of v w.r.t. T.

- *adjacency* tree: To identify a spanning tree T in F, define *id* of T as the minimum id of its nodes. Now we group all the id's of spanning trees in F which are adjacent to v as a red-black tree, and call it adjacency tree of v. So v is linked with T if id of T appears in adjacency tree of v. And this happens only if v has a sub-adjacent tree w.r.t. T.

We will maintain adjacency tree of v even if it has been switched "off." Consequently, adjacency trees should be updated accordingly only if edge updates occur in F. Note that the structure of sub-adjacency tree is similar to that of path tree. Both has nodes sorted according to Euler-tour order.

Updates of adjacency structures to reflect an edge cut is shown briefly here. We renew adjacency structures of vertices in C one by one. Suppose the edge cut occurs on T, splitting T into T_1 and T_2. For a vertex v in C, first check whether T is in adjacency tree of v. If "yes", remove id of T from it, and update related sub-adjacent tree to reflect edge cut of T. For T_j ($j = 1, 2$), add its id into adjacent tree of v if it is adjacent to v (determined by whether a sub-adjacent tree of v exists for T_j). As for every vertex in C, we have to test and do such maintenance when necessary. Therefore, total time is $\widetilde{O}(m^{1/2})$ since $|C|$ is $O(m^{1/2})$. We state it as a lemma.

Lemma 3.3. *Adjacency structures for C w.r.t. F can be renewed in $\widetilde{O}(m^{1/2})$ time subjected to an edge update on F. And given a query of whether a vertex v in C is adjacent to a spanning tree T of G_A, it can be answered in $\widetilde{O}(1)$ time.*

3.3. **The whole structure.** Recall again that the goal is to remove V_A. We will first add some artificial vertices and edges into the subgraph of G induced by $V_B \cup C$, resulting a graph H. (Note that here we include the vertices in $C \setminus S$, rather than just V_C, which is $C \cap S$.) Added vertices or edges are used to restore the "damage" of connectivity due to removal of V_A. After it, we will define a graph G^* based on H.

Recall that components of G_A are classified as either low or high. Details about added artificial vertices or edges are as follows.

- Added by path graphs:
  For every spanning tree T in F, build a path graph of V_B w.r.t. T.

- Added by high components:
  For every high component P in G_A, add a *meta-vertex*. Link all vertices in C adjacent to P with it. Also link the 1st vertex of the path graph of V_B w.r.t. the spanning tree corresponding to P with it.

- Added by low components:
  For every low component Q in G_A, first build a complete graph within the vertices in C adjacent to Q. Next, again link all these vertices with the 1st vertex of the path graph of V_B w.r.t. the spanning tree corresponding to Q.

After these, graph H can be formally defined as:

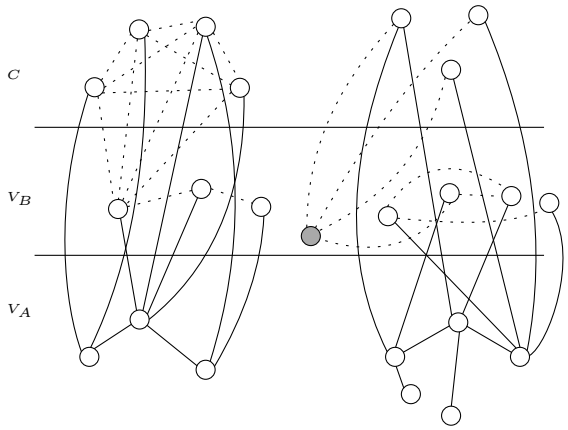

FIGURE 2. An example of the structure. Irrelevant edges within each vertex set are omitted for clarity. *Solid* edges are ones in G, while *dotted* ones denote added artificial edges. *Grey* vertex in the V_B layer indicates a meta-vertex. Note that the left component of V_A is a low component; whereas the right one is high. We build a complete graph within vertices in C w.r.t. the low component.

- Vertex set $V(H)$ of H: $V_B \cup C \cup M$, where M is the set of added meta-vertices. Since the degree of a high component is $> m^{1/4}$, and vertices in $V_B \cup C$ are of degree $> m^{1/4}$, H has at most $O(m^{3/4})$ vertices.
- Edge set $E(H)$ of H: Original edges of G within $V_B \cup C$, and the set of added artificial edges.

See Fig. 2 for an example. H is a multigraph. And in fact, we are only interested in its edge set. So there is a value $D[u, v] > 0$ of edge multiplicity for an edge (u, v) in H. And we say all these $D[u, v] > 0$ constitute an *edge pool*. We will discuss how to maintain it in later part of this subsection.

Now we construct a graph G^*:

- Vertex set $V(G^*)$ of G^*: $V_B \cup V_C \cup M$.
- Edge set $E(G^*)$ of G^*: edge (u, v) with $D[u, v] > 0$, where $u, v \in V(G^*), u \neq v$.

In other words, G^* is a variant of the subgraph of H induced by $V_B \cup V_C \cup M$. It excludes vertices in $C \setminus S$, i.e. only vertices in V_C of C are contained. Also, multiple edges are substituted by single ones. And edges (v, v)'s are removed. The filtration makes the subgraph be simple so that dynamic graph connectivity data structures can be used.

We will *equip* G^* with the worst-case poly-logarithmic update-time structure for dynamic graph connectivity. And as to update and query time of the structure, we re-state a theorem from [24].

Theorem 3.4 ([24])**.** *Given an initial graph with m edges and n vertices, for any constant c and any sequence of updates and queries, the randomized structure supports worst-cast time $O(\lg^4 n)$ per edge insertion, $O(\lg^5 n)$ per edge deletion, and $O(\lg n / \lg \lg n)$ per query. The answer to each query is correct if the answer*

is "yes" and is correct with probability $\geq 1 - 1/n^c$ *if the answer is "no." The preprocessing cost of it is* $O(m \lg^3 n + n \lg^4 n)$.

About edge pool $D[u, v]$'s, a red-black tree suffices to store them, with each entry $D[u, v]$ indexed by $u+nv$ (assuming $u \leq v$). Along the process of updates, we might increment or decrement $D[u, v]$'s. When a $D[u, v]$ decrements to 0, we remove it from edge pool. And if both u and v are vertices in G^* and $u \neq v$, edge (u, v) is deleted from G^*; when a $D[u, v]$ increments to 1, we add it to edge pool. And if both are vertices in G^* and $u \neq v$, edge (u, v) is inserted into G^*.

G^* captures the property of connectivity, which is stated in the following lemma.

Lemma 3.5. *For two vertices* $u, v \in V_B \cup V_C$*, they are connected in the subgraph of* G *induced by* S *iff they are connected in* G^**.*

Proof. G^* is a variant of G without V_A. Connectivity within vertices in V_B through V_A is captured by path graphs built. For V_C, connectivity within it is *either* restored by linking with the same meta-vertex, *or* complete graph built. Last, for connectivity between V_C and V_B, it is restored by the 1st vertex of path graph linked with the meta-vertex, or all relevant vertices in V_C.

Consider a path between u and v in the subgraph induced by S, segments of the path consisting only of vertices in V_A can be eliminated, as "via V_A" connectivity is captured as discussed above. The lemma follows trivially. □

3.4. **Update and query.** In this section, we discuss how vertex updates are reflected efficiently in the data structures. The main task is to keep $D[u, v]$'s being consistent with G. As G^* will select a subset of edges in edge pool, it might also need to be updated. About queries, we will first give its algorithm, then discuss its error probability since G^* is equipped with a randomized data structure.

- **Update**
  - $v \in A$: Suppose v is switched "on." We first restore the relation between v and C, then its relation with vertices in V_B. Edges between vertices of V_A and v are considered in the last step.

    Mark v "active." First restore edges between v and C. For every vertex in C adjacent to v, build a sub-adjacency tree (consisting only of node v) of this vertex of C, and add id of v to its adjacency tree. After this, complete graph among these vertices in C are constructed, as the degree of v is $\leq m^{1/4}$, i.e. a low component. Cost is $\widetilde{O}(m^{1/2})$, dominated by establishing the complete graph.

    Continue to restore the relation of v with vertices in V_B. Build subpath tree, and path tree of v (like sub-adjacency tree, path tree contains only v if existed). Add to edge pool the path-graph edges associated with v ("*add*" means incrementing the involved entry $D[u, v]$), and also the edges between some vertices in C and the 1st vertex of the path graph. Cost is $\widetilde{O}(m^{1/4})$.

    Last, insert edges between vertices in V_A and v one by one. $\widetilde{O}(\sqrt{m})$ deterministic structure handles these in $\widetilde{O}(m^{3/4})$ time. Note that as the degree of v is $\leq m^{1/4}$, linking or cutting among F occurs at most $O(m^{1/4})$ times. According to Lemma 3.2, path graphs can be maintained in $\widetilde{O}(m^{1/4})$ time. And Lemma 3.3 shows that adjacency structures can be updated in $\widetilde{O}(m^{3/4})$ time.

The number of components in G_A affected is at most $O(m^{1/4})$. For every high component, with the help of adjacency structures, edges between C and the meta-vertex can be determined in $\widetilde{O}(m^{1/2})$ time according to Lemma 3.3, since $|C|$ is at most $O(m^{1/2})$; for every low component, as the degree of a low component is $\leq m^{1/4}$, $\widetilde{O}(m^{1/2})$ suffices to build the complete graph for vertices in C adjacent to the low component. In a word, no matter which case, the time for each component is bounded by $\widetilde{O}(m^{1/2})$. So cost for this step is $\widetilde{O}(m^{3/4})$.
The case of v "off" is just a reverse process of above. In summary, for an update of $v \in A$, cost is bounded by $\widetilde{O}(m^{3/4})$.

- $v \in B$: For variety, suppose v is "off." First destroy the relation of v with vertices in V_A. Lemma 3.2 provides a bound of $\widetilde{O}(m^{1/2})$ for updating path graphs.
  Note that v might be the 1st vertex of some path graphs before this. So we need to handle it. v can be adjacent to at most $m^{1/2}$ components in G_A, as its degree is $\leq m^{1/2}$. For a high component, it is trivial, as only one edge linking it with a meta-vertex; for a low component, since only $\leq m^{1/4}$ edges can be outward for a low component, $\widetilde{O}(m^{1/4})$ suffices for updating edges between the vertices in C adjacent to the low component and the 1st vertex. In a word, cost is $\widetilde{O}(m^{3/4})$.
  After this, all remaining edges incident on v are edges in G. Artificial edges are all removed. For the next step, just delete these edges linking v with vertices in V_B and vertices in C. Cost is $\widetilde{O}(m^{1/2})$. Finally, mark v "inactive."
  In total, for this case cost is bounded by $\widetilde{O}(m^{3/4})$. The other case is also a reverse procedure. Therefore, update cost for this type is $\widetilde{O}(m^{3/4})$.

- $v \in C$: As there are only $O(m^{3/4})$ vertices in G^*, update time is $\widetilde{O}(m^{3/4})$. Recall that relevant $D[u, v]$'s are left intact, and we do not destroy (if v is "off") adjacency structure of v. Mark v as necessary. Cost is $\widetilde{O}(m^{3/4})$.

- **Query**
  - Given two active vertices u and v, our first attempt is to express them as equivalent vertices in G^* (in the sense of being connected). Recall that vertices in G^* are $V_B \cup V_C \cup M$, where M are the set of meta-vertices induced by G_A.
    It is done as follows: If both are already in $V_B \cup V_C$, it is solved; otherwise, for any one not in, if it is in a high component, just replace it with the meta-vertex; if it is in a low component, exhaustively search outward edges of the low component for a vertex of G^*.
    Connectivity within G^* can be answered by the randomized dynamic connectivity structure; whereas for other cases, u and v are connected iff they are in the same component of G_A, of which queries can be answered by the deterministic dynamic connectivity structure for G_A.

Cost mainly comes from exhaustive search incurred by a low component. So query time is $\widetilde{O}(m^{1/4})$.
About its correctness, if u and v are both vertices in $V_B \cup V_C$, it follows from Lemma 3.5; otherwise, for any one not in, we only replace it with an equivalent vertex of G^*. If such an equivalent vertex cannot be found, the query vertex is on an "island" of G_A. And for this case, they are connected iff they are on the same island.

Lastly, we discuss the error probability of query operations. Note that a deterministic structure is adopted for G_A. Therefore, topology of G^* is guaranteed to be correct. As to answers to queries, they are *either* answered by the deterministic structure if there exists a query vertex being on an "island" of G_A, *or* answered by the randomized structure if both query vertices are (replaced with) vertices in G^*. In the former case, correctness is always ensured; whereas the latter one might answer erroneously.

We claim that error probability of the latter is still bounded by $1 - 1/n^c$, being in consistent with Theorem 3.4. As discussed above, topology of G^* is correct. And the approach in [24] maintains a private *witness* of spanning forest of G^*. And the property is that after every update, the witness is a spanning forest of G^* with probability $\geq 1 - 1/n^c$, which is shown in Theorem 3.4. It is this property that ensures answers are correct w.h.p. And for our data structure, after every vertex update (which is transformed to a sequence of edge updates in G^*), the witness for G^* is a spanning forest of G^* w.h.p. after the update, as we can just focus on the correctness of the witness at the point after the last edge update transformed. Consequently error probability is negligible.

3.5. **Space complexity and preprocessing time.** In this section, we show that the data structure occupies linear space. The space usage consists of deterministic connectivity structure on G_A, edge pool, randomized connectivity structure on G^*, and some other assistant structures, e.g. adjacency structures. Later, we discuss the preprocessing time. There, it can be observed that the cost mainly comes from $D[u, v]$ entries constructed for complete-graph edges within C.

- **Space complexity**
  Deterministic connectivity structure for G_A occupies space proportional to size of G_A. And edges in G_A are all original edges in G. So this part uses linear space.

  For the $D[u, v]$'s, entries corresponding to original edges in G conforms to linear-space usage. For artificial edges introduced by path graphs, note that every original edge in G between a vertex of V_B and one of V_A is replaced by at most 2 edges within path graphs. For artificial edges connecting the 1st vertex of a path graph with some vertices of C, each one of them is brought because there is an original edge linking a vertex of C with a low component in G_A. Edges between vertices of C and meta-vertices can be accounted for in the same way. Until now, every original edge in G is attributed $O(1)$ times. So space usage excluding complete graphs built is linear. Last for excluded edges constructed for complete graphs, as only multiplicities are recorded, and $|C| = O(m^{1/2})$, $O(m)$ space suffices. In summary, $D[u, v]$'s size is linear.

For randomized connectivity structure on G^*, its space usage is $O(m' + n' \lg^3 n')$, for a graph with m' edges, n' vertices. (The $\lg^3 n'$ factor is due to $O(\lg n')$ *cutset* data structures, which are represented by ET-trees maintaining $O(\lg^2 n')$ values. See [24] for details.) Edge set of G^* is a subset of $D[u, v]$'s, therefore its cardinality is $O(m)$. And as to n', its value for G^* is $O(m^{3/4})$. In summary, space usage of randomized structure is linear.

For other parts like adjacency structures, every appearance of a vertex of V_A in an sub-adjacency tree corresponds to an original edge in G between C and V_A. Thus it is linear. In summary, combining all of these, we can see that total space complexity is linear.

- **Preprocessing time**

  As to preprocessing time, deterministic structure and randomized one both cost $\widetilde{O}(m)$ time. Similar to space usage, excluding complete graphs, every edge between V_A and $V_B \cup C$ induces $O(1)$ artificial edges, and $\widetilde{O}(1)$ preprocessing cost suffices.

  For the uncovered complete graphs, note that for every edge from a vertex of C to a low component, $O(m^{1/4})$ edges are added in complete graphs (the degree of a low component is $\leq m^{1/4}$). So, preprocessing time for complete graphs is bounded by $\widetilde{O}(m^{5/4})$. The total preprocessing time is dominated by this part, and thus it is $\widetilde{O}(m^{5/4})$.

4. Extension to edge updates

The structure described in Section 3 can be extended to support edge updates, although the resulted edge-update time is huge, costing the same as vertex update. Also remember that we work on a variant of G, which is mentioned in the beginning of Section 3. And this variant has the property of $m = \Omega(n)$ during its lifetime.

First note that, to insert an edge (u, v), the simple trick of creating a new vertex z adjacent to only u and v, and then inserting z would not work here [5, 6]. The point is that this might violate the type of u (or v), i.e., which set (A, B or C) it belongs to. Meanwhile the property of a component of G_A being low or high might also be hurt.

We circumvent these troubles by the following steps.

Step 1 Turn u (or v) "off" if u (or v) is not "off" currently.

Step 2 Update edge set E.

Step 3 Change the type(s) of u (and v). Type transition between C and B is the only case requiring care. For other cases, just mark the new type(s), e.g., if a former A vertex v becomes a B type, just mark type of v as B. W.l.o.g., suppose the degree of a former B vertex u increases to $> m^{1/2}$, i.e. a new C vertex. Note that $\deg_G(u)$ is just $\Theta(m^{1/2})$, as it is just above the threshold. First, construct its adjacency structure, which can be accomplished in $\widetilde{O}(m^{1/2})$ time; Second, insert $D[u, v]$ entries into edge pool. As $\deg_G(u)$ is $\Theta(m^{1/2})$, the worst case is that it is adjacent to $\Theta(m^{1/2})$ distinct low components. $\widetilde{O}(m^{3/4})$ suffices for calculating $D[u, v]$'s.

Step 4 Restore u (or v) to its former active state if it is "on" before the edge update.

Total cost of handling an edge update is $\widetilde{O}(m^{3/4})$. Although the approach above seems like overkill, note that in the worst case, it cannot be done better. As type transitions should be tackled, and in the worst case, for example, a V_A vertex become V_B, then it should be deleted from V_A via turning it "off" first, and this could cost as much as $\widetilde{O}(m^{3/4})$. As a comparison, in [6], by adopting the simple trick of inserting a vertex z, the degree is only 2, however amortized cost for edge update is still $\widetilde{O}(m^{2/3})$, without any advantage over vertex update.

The last thing is about the value of m. As m is used as a parameter for division of vertices and partition of components, m should not deviate from the true number of edges m' currently in graph G by too much. So when detecting $|m' - m|$ first $\geq \frac{1}{2}m^{1/2}$, we make a copy G' of the current graph G, and partition vertices based on the value of m'. During the next $\frac{1}{4}m^{1/2}$ updates, we preprocess G' in $\widetilde{O}((m')^{5/4})$ time, which is $\widetilde{O}(m^{5/4})$ as $m' = \Theta(m)$. The preprocessing is done in the background, and is evenly distributed over the $\frac{1}{4}m^{1/2}$ updates. After preprocessing, we start catching-up process of performing $\frac{1}{2}m^{1/2}$ updates during the next followed $\frac{1}{4}m^{1/2}$ updates of G. By the end of $\frac{1}{2}m^{1/2}$ updates of G, we replace G with G'.

For the past $\frac{1}{2}m^{1/2}$ updates (supported by G), each only spends $\widetilde{O}(m^{3/4})$ worst-case update time on the preprocessing and catching-up of G'. So worst-case update time is still $\widetilde{O}(m^{3/4})$. And the crucial point is that for all past updates, using m is just fine, as difference between m and true value is $\leq m^{1/2}$, i.e. asymptotically these two values are the same.

Now we replace G with G'. Suppose the current number of edges is m''. And we face a similar circumstance of $|m'' - m'| \leq \frac{1}{2}(m')^{1/2}$ to the former G. Just begin the process above again, if at some point detecting $|m'' - m'|$ first $\geq \frac{1}{2}(m')^{1/2}$. Summing up the paper, we have shown Theorem 1.1.

5. Conclusions

In this paper, we improve worst-case update time of dynamic subgraph connectivity to $\widetilde{O}(m^{3/4})$. For further research, one obvious question is, for amortized running time, whether $\widetilde{O}(m^{2/3})$ can be improved. As discussed in [5, 6], detecting triangles in digraphs can be reduced to dynamic subgraph connectivity with $O(m)$ offline vertex updates. Therefore an update time better than $\sqrt{m}$ appears hard without FMM, as the best running time for finding triangles without FMM is $O(m^{3/2})$ [23].

Another intriguing question is whether the product of update and query time can be sublinear, i.e. $o(m)$. Current technique deployed is removal of vertices, and then construct a graph G^*, of which the number of vertices is limited. The idea was used in nearly every algorithm proposed [5, 6, 10]. It has a drawback, however. Assume the update time is $O(m^{1-\epsilon})$, for $\epsilon > 0$, then for vertices with degree $> m^{1-\epsilon}$, the switch of state for each one of them cannot be explicitly notified to all their neighbors, as update time should be $O(m^{1-\epsilon})$. So others can be *blind* to states of these high-degree ones. Therefore explicit checks need to be done, whereas the number of high-degree vertices can be $\Omega(m^{\epsilon})$. So, for this question, a new approach might be needed. And solving this problem seems to be hard, as indicated in [18]. There, they essentially showed that OMv (*Online Boolean Matrix-Vector Multiplication*) conjecture (see [18] for details) rules out algorithms with *polynomial* preprocessing time, and the product of amortized update and query

time being $o(m)$. Nonetheless, there are many other open problems in this area. See [6] for details.

References

[1] Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 434–443. IEEE, 2014.

[2] Surender Baswana, Shreejit Ray Chaudhury, Keerti Choudhary, and Shahbaz Khan. Dynamic dfs in undirected graphs: breaking the o (m) barrier. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 730–739. SIAM, 2016.

[3] Aaron Bernstein and David Karger. Improved distance sensitivity oracles via random sampling. In *Proceedings of the nineteenth annual ACM-SIAM Symposium on Discrete algorithms*, pages 34–43. Society for Industrial and Applied Mathematics, 2008.

[4] Aaron Bernstein and David Karger. A nearly optimal oracle for avoiding failed vertices and edges. In *Proceedings of the forty-first annual ACM Symposium on Theory of computing*, pages 101–110. ACM, 2009.

[5] Timothy M. Chan. Dynamic subgraph connectivity with geometric applications. In *Proceedings of the thiry-fourth annual ACM Symposium on Theory of Computing*, pages 7–13. ACM, 2002.

[6] Timothy M. Chan, Mihai Pătraşcu, and Liam Roditty. Dynamic connectivity: Connecting to networks and geometry. *SIAM Journal on Computing*, 40(2):333–349, 2011.

[7] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251 – 280, 1990. Computational algebraic complexity editorial.

[8] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.

[9] Camil Demetrescu, Mikkel Thorup, Rezaul Alam Chowdhury, and Vijaya Ramachandran. Oracles for distances avoiding a failed node or link. *SIAM Journal on Computing*, 37(5):1299–1318, 2008.

[10] Ran Duan. New data structures for subgraph connectivity. In *Automata, Languages and Programming*, pages 201–212. Springer, 2010.

[11] Ran Duan and Seth Pettie. Dual-failure distance and connectivity oracles. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 506–515. Society for Industrial and Applied Mathematics, 2009.

[12] Ran Duan and Seth Pettie. Connectivity oracles for failure prone graphs. In *Proceedings of the forty-second ACM Symposium on Theory of Computing*, pages 465–474. ACM, 2010.

[13] Ran Duan and Seth Pettie. Connectivity oracles for graphs subject to vertex failures. *CoRR*, abs/1607.06865, 2016.

[14] David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Amnon Nissenzweig. Sparsification-a technique for speeding up dynamic graph algorithms. *Journal of the ACM*, 44(5):669–696, 1997.

[15] Greg N. Frederickson. Data structures for on-line updating of minimum spanning trees, with applications. *SIAM Journal on Computing*, 14(4):781–798, 1985.

[16] Daniele Frigioni and Giuseppe F. Italiano. Dynamically switching vertices in planar graphs. *Algorithmica*, 28(1):76–103, 2000.

[17] David Gibb, Bruce M. Kapron, Valerie King, and Nolan Thorn. Dynamic graph connectivity with improved worst case update time and sublinear space. *CoRR*, abs/1509.06464, 2015.

[18] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proceedings of the forty-seventh Annual ACM on Symposium on Theory of Computing*, pages 21–30. ACM, 2015.

[19] Monika Henzinger and Stefan Neumann. Incremental and fully dynamic subgraph connectivity for emergency planning. *arXiv preprint arXiv:1611.05248*, 2016.

[20] Monika R. Henzinger and Valerie King. Randomized fully dynamic graph algorithms with polylogarithmic time per operation. *Journal of the ACM*, 46(4):502–516, 1999.

[21] Jacob Holm, Kristian De Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *Journal of the ACM*, 48(4):723–760, 2001.
[22] Shang-En Huang, Dawei Huang, Tsvi Kopelowitz, and Seth Pettie. Fully dynamic connectivity in o(log n(log log n)^2) amortized expected time. *CoRR*, abs/1609.05867, 2016.
[23] Alon Itai and Michael Rodeh. Finding a minimum circuit in a graph. *SIAM Journal on Computing*, 7(4):413–423, 1978.
[24] Bruce M. Kapron, Valerie King, and Ben Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In *Proceedings of the twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1131–1142. SIAM, 2013.
[25] Casper Kejlberg-Rasmussen, Tsvi Kopelowitz, Seth Pettie, and Mikkel Thorup. Faster worst case deterministic dynamic connectivity. *CoRR*, abs/1507.05944, 2015.
[26] Peter Bro Miltersen, Sairam Subramanian, Jeffrey Scott Vitter, and Roberto Tamassia. Complexity models for incremental computation. *Theoretical Computer Science*, 130(1):203–236, 1994.
[27] Mihai Pătraşcu and Erik D. Demaine. Logarithmic lower bounds in the cell-probe model. *SIAM Journal on Computing*, 35(4):932–963, 2006. Special issue of selected papers from the 36th ACM Symposium on Theory of Computing, 2004.
[28] Mihai Pătraşcu and Mikkel Thorup. Planning for fast connectivity updates. In *Proceedings of the forty-eighth IEEE Symposium on Foundations of Computer Science*, pages 263–271. IEEE, 2007.
[29] Robert E. Tarjan. Dynamic trees as search trees via Euler tours, applied to the network simplex algorithm. *Mathematical Programming*, 78(2):169–177, 1997.
[30] Mikkel Thorup. Near-optimal fully-dynamic graph connectivity. In *Proceedings of the thirty-second annual ACM Symposium on Theory of Computing*, pages 343–350. ACM, 2000.
[31] Zhengyu Wang. An improved randomized data structure for dynamic graph connectivity. *CoRR*, abs/1510.04590, 2015.
[32] Christian Wulff-Nilsen. Faster deterministic fully-dynamic graph connectivity. In *Proceedings of the twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1757–1769. SIAM, 2013.

Appendix A. Extension to $\widetilde{O}(m^{3/4+\epsilon})$ update and $\widetilde{O}(m^{1/4-\epsilon})$ query time

Here we show the scaffolding of the structure, and discuss how to adjust parameters to achieve a query time of $\widetilde{O}(m^{1/4-\epsilon})$. First partition vertices according to degrees.

- C: Vertices with $\deg_G(v) > m^{1-b}$
- B: Vertices with $m^{1-a} < \deg_G(v) \leq m^{1-b}$
- A: Vertices with $\deg_G(v) \leq m^{1-a}$

And a component of G_A is high if its total degree is $> m^{1-c}$. Otherwise, it is low.

As a low component can be of degree as high as m^{1-c}, query time should be $\widetilde{O}(m^{1-c})$. Consequently, we can anticipate an update time of $\widetilde{O}(m^c)$. (Section 5 explains why we expect this.)

To achieve the desired update time, the following inequalities should be satisfied.

- Update in A:
  (1) Maintaining a spanning forest of G_A: $\sqrt{m} \cdot m^{1-a} \leq m^c$.
  (2) Maintaining complete graphs w.r.t. low components: $m^{1-a} \cdot m^{2(1-c)} \leq m^c$.
  (3) Maintaining edges between C and meta-vertices: $m^{1-a} \cdot m^b \leq m^c$.

- Update of a B vertex v:
  (4) Dominant update cost comes from v being the 1st vertex of path graphs: $m^{1-b} \cdot m^{1-c} \leq m^c$.

- Update in C:
  (5) The number of vertices in G^* should be bounded: $m^a \leq m^c$.

From (2) and (5), we have $3 \leq a + 3c \leq 4c$. So now you could understand why the structure can at most achieve an update time of no better than $\widetilde{O}(m^{3/4})$, as c should be $\geq 3/4$. Meanwhile, it tells us c could be $3/4 + \epsilon$. By setting b to be $1/2$ and a to be $3/4$, all inequalities are again satisfied. So now we have a structure with $\widetilde{O}(m^{3/4+\epsilon})$ update and $\widetilde{O}(m^{1/4-\epsilon})$ query time.

Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China
E-mail address: duanran@mail.tsinghua.edu.cn

Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China
E-mail address: le-zhang12@mails.tsinghua.edu.cn